

(19)



Europäisches Patentamt  
European Patent Office  
Office européen des brevets



(11) Publication number:

**0 413 537 A2**

(12)

## EUROPEAN PATENT APPLICATION

(21) Application number: **90308878.9**

(51) Int. Cl.<sup>5</sup>: **H04L 12/54**

(22) Date of filing: **13.08.90**

The title of the invention has been amended  
(Guidelines for Examination in the EPO, A-III,  
7.3).

(30) Priority: **15.08.89 GB 8918553**

(43) Date of publication of application:  
**20.02.91 Bulletin 91/08**

(84) Designated Contracting States:  
**DE ES FR GB IT**

(71) Applicant: **DIGITAL EQUIPMENT  
INTERNATIONAL LIMITED  
1 Grand Places  
CH-1700 Fribourg(CH)**

(72) Inventor: **Scannell, Niamh  
1200 Dale Avenue No.17  
Mountain View, California 94040(US)  
Inventor: Redmond, Anthony, John  
58 Fairways  
Rathfarnham, Dublin 14(IR)  
Inventor: Himbaut, Serge Elvina Hills C2  
1187 Route de St. Jean  
F-06600 Antibes(FR)  
Inventor: Bares, Pascale  
13 Lantanas Hameaux du Soleil  
F-06270 Villeneuve Loubet(FR)  
Inventor: Clark, Alison  
52 Foxhill Road  
Reading Berkshire(GB)  
Inventor: Dawson, Stuart Douglas  
40 Lorne Street  
Reading Berkshire(GB)**

(74) Representative: **Goodman, Christopher et al  
Eric Potter & Clarkson St. Mary's Court St.  
Mary's Gate  
Nottingham NG1 1LE(GB)**

(54) **Electronic mail message control system.**

(57) An electronic mail system in which in a processing station messages are sorted according to a set of criteria established for particular users to allocate messages and to preferably order the sequence in which messages are presented to each user.

**EP 0 413 537 A2**

## MESSAGE CONTROL SYSTEM

The present invention relates to message control systems.

Digital communication systems of the "electronic message" or "electronic mail" type are well established. In such systems, several (often a very large number) of work stations are interconnected by a system which allows users at the work stations to send messages to each other. Such messages are the electronic equivalent of letters and memoranda.

In such a system, there will normally be a standard format defined for such messages (or, often, a number of slightly different formats for different purposes). The format will normally have several fields, including a sender field, a date field, an addressee field, and a message body field; in addition, there will often be further fields, including in particular a copy-to field and a subject field. The addressee and copy-to fields will often each be capable of holding more than one identifier, so that a message can be sent to several addressees and copied to several more. (The addressee and copy-to fields are normally functionally equivalent as far as the system is concerned, with the distinction between the two being of significance only to the users.)

In such a system, when a user generates a message, the system will automatically transmit it to the addressees (including the copy-tos). At each of the stations to which it is sent, it will be stored in an (electronic) folder (which may well be termed the "In-tray" folder). When the user at that station chooses to inspect the contents of the In-tray, they will extract the various messages from it one after another, inspect each, and take whatever action they think appropriate, such as deleting the message, printing it out, or storing it in some other folder.

In practice, it is found that such a procedure can become burdensome, because the number of messages which accumulate in the In-tray can become so large that a considerable amount of time and effort is required for the user to work through the accumulation of messages in the In-tray.

The present invention is to provide an electronic mail system comprising a plurality of stations between which messages can be sent, and including classification means in each station which assigns different priorities to the messages in dependence on matching of each message to a set of criteria defined by a set of rules. Each station preferably includes a message store for accumulating incoming messages and in which the classification means are adapted to operate on the messages so accumulated in the message store to

order the sequence in which the messages are presented to the user without changing the messages themselves.

The invention also provides a message processing station for an electronic mail system comprising a message store for storing incoming messages and having a rules store containing a plurality of rules relating to the status of received messages. The station also preferably includes a rules test circuit coupled to the message store and to the rules store and including comparison means for comparing a message stored in the message store with the rules to achieve a set of comparison results; and routing means for routing each message to a location selected by the set of comparison results for that message.

The present invention further provides a method of controlling messages in an electronic mail system comprising a plurality of stations between which messages can be sent, the method including receiving and storing an incoming message and comparing the stored incoming message with a plurality of rules to produce a set of comparison results for that message. The method also preferably allocates at least one location for subsequent storage of the message in accordance with the comparison results for that message.

A message sorting and processing apparatus embodying the invention will now be described, by way of example, with reference to the drawings, in which:

Fig. 1 is a general block diagram of the system; and

Fig. 2 is more detailed block diagram of the message and rules stores.

### Organization - general

Fig. 1 shows in simplified form those parts of a work station which form the present apparatus. It will be understood that the work station will include other parts, not shown, providing general data processing facilities. Also, some or all of the functions of the work station may in fact be implemented by means of a host computer shared by a number of users and which presents to each user the facilities of an independent work station under the sole control of that user.

Referring to Fig. 1, there is an I/O port 10 by means of which messages can be transmitted to and received from other work stations, via a message transmission network (which may include or consist of a host computer). Incoming messages pass through the I/O port to a message store 11,

which holds them in readiness for processing by the rest of the apparatus. As will be explained later, the apparatus operates in accordance with a number of rules, and these are stored in a rules store 12. The messages are matched against the rules in a rules test unit 13, which controls a message switch 14. The messages stored in store 11 are passed through the message switch back to the I/O port 10 for transmission to other work stations, and on to a main folder store 15; the folders include an In-tray folder 16 and other folders Folder-1 17, Folder-2 18, etc.

Information about the messages is also stored in a message log unit 19. A keyboard 20 enables the user to control the operation of the apparatus, and is coupled to a monitor 21 on which various operations, and the contents of the message log unit 19, can be displayed. A rules control unit 22 is coupled to the keyboard 20 and the rules store 12. This unit enables the user to generate new rules and modify existing rules; existing rules are extracted from the rules store 12 for this purpose, and new or modified rules are returned to that store.

### Operating modes

The apparatus may be operated in either continuous or batch mode, under control of a general control unit 23. This unit controls the remaining units of the apparatus, and also controls the mode of operation of the apparatus. The mode is selected by means of the keyboard 20.

In continuous mode, each message is processed as it is received. This involves applying all the rules in turn to the message, and taking the appropriate actions. In the batch mode, incoming messages are accumulated in the message store 11, and processed at selected times; the processing times may be preprogrammed into the system (e.g. at times 0845 and 1245, so that the processed messages are available for the user to deal with when they come in in the morning and back from lunch), or when demanded by the user.

In the case where the apparatus operates on demand from the user, it initially presents the user with an indication of the number of messages in the message store 11. The user can then select whether the apparatus shall operate in foreground mode (so that nothing else can be done until the message processing has been completed) or in background mode (so that the user can proceed with some other task, with the apparatus performing the message processing in the intervals in which that other task is not utilizing the resources of the work station).

The modes of operation just mentioned are

both implicitly automatic. The apparatus can also be operated in a manual mode. For this, the apparatus presents the user with the proposed actions for each message, but awaits confirmation from them before performing those actions. For this, the apparatus preferably presents the proposed actions on the screen 21, so that the user can select the proposed actions one after another, e.g. by means of a cursor, for performance or not. It may be desirable for the screen to be split or windowed, so that the message can also be seen; the message is preferably scrollable on its part of the screen for this.

For this mode, the user may set the apparatus to process the messages automatically for the current batch (i.e. the current contents of the message store 11), but to revert to the manual mode thereafter.

If the user decides not to use the present apparatus, then the messages in the message store 11 are simply transferred directly to the In-tray folder 16. The user can then obtain an indexed listing of them, read them, and otherwise process them in conventional manner.

There is a further mode, intermediate between those described above, in which the apparatus operates but the only action taken is to assign priorities to the messages as they are passed from the message store 11 to the In-tray folder 16.

Means are also preferably provided whereby the user can pass a selected message from the In-tray to the rules test unit 13. This enables the user to determine what actions the rules recommend on a message which has reached the In-tray without those actions having already been performed.

To forward a message to another user, it is convenient to utilize a technique in which, from a formal point of view, a new message is constructed. This involves generating appropriate fields for the new message, the new addressee field indicating the user to which the message is to be forwarded and the new sender field indicating the user from whom it is being forwarded. This forwarding technique is known per se. The actual construction of the new message is performed by the message switch 14, which also performs other data processing functions such as appending priority levels to messages being passed to the In-tray folder 16.

The contents of the new subject field may depend on the mode of operation; if the mode is automatic, then the contents of the new subject field may be copied from the subject field of the original message; if the mode involves monitoring or participation by the user, then the system may offer the user the option of either inserting a new subject or having the contents of the old subject field copied. The body of the new message is constructed by concatenating all the fields of the

old message.

The messages to be forwarded on to other users are passed from the message switch 14 to an output message store 24, from which they are passed to the I/O port 10.

Statistical information relating to the operation of the apparatus, e.g. the total number of messages and the number of messages forwarded, is accumulated in the message log unit 19. This information can be extracted by the users if they wish to monitor and review the operation of the apparatus.

When message processing has been completed, by whatever variation of mode, the processed messages will be in the In-tray folder 16, in prioritized order for the user to attend to. (Of course, the users can choose to transfer the messages direct to the In-tray, without processing, if they desire.)

### Organization - details

Fig. 2 shows the organization of the message store 11, the rules store 12, and the rules test unit 13. The message store 11 contains a plurality of message storage units 25. Each new incoming message is stored in the next free unit 25, and when the apparatus runs to evaluate messages, each message in turn is evaluated and removed from its unit, leaving that unit free. (Two pointers are used in a conventional manner to keep track of which units are occupied by messages.) Similarly, the rules store 12 contains a plurality of rules storage units 35.

Considering now the structure of a message storage unit 25, this is organized to match the structure of the messages. Accordingly it consists of several sub-units or fields, corresponding to the fields of the messages. Specifically, it has the following fields:

- a sender field 26 which contains the identity of the sender. (This identity is a unique binary identifier used by the system. Depending on the specific design of the system, the identifier may be chosen in a form which also has a mnemonic significance, or the system may include conversion tables which convert the identifier to a human-recognizable name.)
- a date field 27. This is not utilized in the present apparatus.
- an addressee field 28. This contains the identities of the addressees to which the message is being sent, in the form of identifiers like the sender identifier. There can be more than one addressee.
- a copy-to field 29. This contains the identities of the addressees to which the message is being copied. As noted above, the distinction between

addressees and copy-to is of relevance only to the users of the system, not to the system itself. A copy-to addressee is usually sent the message merely for information, without any requirement that they should take any particular action.

- a subject field 30. This will contain a short heading identifying, to the addressees, the subject matter of the message.

- a body field 31. This will contain the main text of the message, and will be of highly variable length.

In practice, the storage units 25 and the fields in each will usually be implemented by assigning suitable regions in a general-purpose memory unit. This reduces the cost of the apparatus and allows memory space to be utilized effectively. Techniques for achieving this are well established.

The apparatus contains a number of rules in the rules store 12. Each rule consists of two parts, a tests or conditions part and an actions part. The apparatus operates by testing a message against each rule in turn. If the message matches the conditions part of the rule, then the apparatus takes the actions prescribed by the actions part of the same rule.

Considering now the structure of the conditions part of a rule, this matches the structure of the messages, but with significant differences. This conditions part structure can be explained most easily by explaining the functions which it performs or satisfies.

Broadly, the user can set up a number of rules. Each rule tests whether the messages satisfy certain conditions concerning who the sender is, who the addressees and/or copy-to are and their numbers, and the nature of the subject and a definable initial part of the message body field. If a rule is satisfied, then the actions which result can be assigning a priority level to the message, filing it in one or more selected files, and forwarding it to further addressees.

Considering now the structure of a rule storage unit 35, this is divided into two parts, a tests part 35A and an actions part 35B. The tests part 35A is organized to match the structure of the messages, but with significant differences. The sub-units or fields of the tests part 35A are as follows:

- a rule name field 36. This contains a name chosen by the user when the rule is being constructed, to aid them in recognizing the nature of the rule if they want to review or amend it later.
- a sender field 37 which contains the identities of equivalent senders. A user may want to treat messages from different senders (e.g. either of two joint chairmen of a committee) in the same way.
- an addressee number limits field 38. (This has no direct counterpart in the message structure.) This contains three parts, each of which can be set to define lower and upper bounds on the number of

addressees; the three parts are for direct addressees, copy-to's, and total addressees (direct plus copy-to's). A user may wish to treat a message directed to a large number of addressees as of low importance, or more generally to treat a message differently depending on the number of direct and/or copy-to addressees.

- an addressees field 39. This contains a list of equivalent addressees. A user may wish to treat a message in different ways depending on who else it is addressed and/or copied to. Since the user may wish to treat direct addressees and copy-to's differently, each addressee in the list is tagged to indicate whether it is to be matched against direct addressees, copy-to's, or both.

- a keyphrase field 40. This contains a list of keyphrases which are to be matched against the contents of the message. Each keyphrase is tagged to indicate whether it is to be matched against subject field, message body fields, or both. The nature of the keyphrases and their matching is discussed in more detail later.

- a keyphrase zone limit field 41. This sets how much of the message body is used for the keyphrase matching. A user may want to look for a keyphrase anywhere in the entire message body, or they may prefer to limit the search for a keyphrase to the initial part of the message body.

The sub-units or fields of the actions part 35B of the rule storage unit 35 are as follows:

- a priority field 45. If the message matches the rule conditions, then it is given the priority level set by this field, which can have a value of between 1 (highest priority) and 5 (lowest priority).

- a file-to field 46. This contains a list of folders in the user's main folder store 15. If the message matches the rule conditions, then it is filed in the appropriate folders.

- a forward-to field 47. This contains a list of addressees; if the message matches the rule conditions, then it is forwarded to these addressees.

The work station may have a distribution list facility, by means of which the user can, when generating a message, enter the distribution list as a direct and/or copy-to addressee. If this facility exists, the apparatus may allow distribution lists to be used in the sender and addressees fields; such a distribution list will be treated as an indirect addressing of its list of addressees.

The rules storage units 35 in the rules store 12 are preferably subdivided into sets so that a plurality of different sets of rules can be stored, any one of which can be selected by the user. This will allow two different users to share the same work station, or a single user to use different rule sets at different times.

In practice, the storage units 25 and 35 and the fields in each will usually be implemented by as-

signing suitable regions in a general-purpose memory unit. This reduces the cost of the apparatus and allows memory space to be utilized effectively. This can be achieved by conventional techniques.

## Operation - details

The message storage units 25 and the rules storage units 35 are coupled to the rules test unit 13 by coupling means, illustrated diagrammatically as 50 and 51, which select the various fields in succession. Corresponding fields are fed to a comparator unit 52 which makes comparisons of successive pairs of fields, and which is coupled to a latch unit 53 which stores the results (both intermediate and final) of the comparisons.

More specifically, comparator 52 first matches the sender field 26 with the addresses in the sender-list field 37. If there is no match, at this or any later stage, then the comparison with that rule is immediately terminated, and the next rule is selected. (Obviously the order of the various tests can be chosen for maximum efficiency; thus it may be preferable to perform this test after the next-described test. The order in which the tests are performed can even be made dependent on the contents of the fields, so that order in which the sender and addressee fields are tested depends on how many comparisons have to be made for each.)

If there is a match, then the comparator next counts the numbers of addresses in the addressees and copy-to's fields 28 and 29 and matches these with the limits set in the addressee limits field 38. The addresses in the fields 28 and 29 are then matched against those in the addressees list field 39.

If the matching so far has been successful, then the subject and body fields 30 and 31 are matched against the keyphrases field 40, under the control of the body-limit field 41. The keyphrases field may contain a number of keyphrases, each of which is made up of keywords which are treated as character strings for matching purposes. The keywords may be combined in logical combinations in the keyphrases. The comparator 52 first matches the individual keywords, and then evaluates the logical combinations to determine whether the keyphrase is satisfied.

Some of the fields in a rule may be left empty; such fields are in general treated as automatic matches. For example, if the sender-list field 37 is left empty, that field will be treated as matching all senders. Thus a rule can operate only on senders, only on keyphrase, or on both in combination.

If a rule is successfully matched, then the latches 53 are set to perform the actions stored in the actions part 35B of the rule storage unit 35, and

the actions are copied into an actions storage unit 54 in the rules test unit 13. When the message has been matched against all rules of the rule set, the actions stored in the actions storage unit 54 are matched against each other, to eliminate duplications of file-tos and forward-tos and to select the highest of whatever priorities have been assigned to the message. The actions are then carried out; the message is filed in the In-tray with the appropriate priority appended to it and in any further appropriate folders and forwarded to the appropriate addressees (if any). The next message, if any, is then selected for processing.

If two or more rules match a message and their actions are different, this may be regarded as indicating an inconsistency between the rules. It would be possible for the rules control unit 22 to test the rules against each other to detect such potential inconsistencies and require the user to correct one or more of the rules to eliminate such inconsistencies. However, such inconsistencies will often be harmless, and it may be preferable to accept some - e.g. forwarding of a message which happens to match the keyphrases fields of two different rules to two different sets of forward-tos. The message log unit may be arranged to record such inconsistencies, so that the user can review them later and decide whether any of them calls for amendment of the rules.

### Rules construction

The rules store 12 is coupled to the rules control unit 22, which is in turn coupled to the keyboard 20 and the monitor 21, as noted above. The operator can copy the rules from the rule store 12 to the rules control unit 22 for inspection and possible amendment, and return them (or newly created rules) to the rules store 12. A rule can also effectively be copied, to serve as a template for the creation of a new but similar rule, by selecting a rule from the rule store and changing its name.

The monitor 21 displays a rule so selected with its fields suitably laid out and labelled, in a manner broadly matching the layout shown in Fig. 2 for the rule storage unit 35. Thus the sender-list field, for example, can conveniently be shown near the top of the screen, with a label such as "Sender List", and with the possible senders to which the rule applies shown in successive lines below the label. The display is preferably scrollable.

The rules control unit incorporates consistency checking means, which check for inconsistencies and require the user to correct them before the rule can be stored in the rules store 12. Such inconsistencies include such matters as duplicated addressees, an upper limit on the number of ad-

ressees which is smaller than the lower limit, and inconsistencies between the limits for direct addressees, indirect addressees, and total addressees.

The keyphrase field is preferably laid out as a number of lines which are effectively ORed; that is, the keyphrase match is satisfied if the conditions of any line are satisfied. On a given line, the keywords are preferably combined into a keyphrase by means of a + symbol for AND and a - symbol for NOT. These conventions enable users to make effective use of the permitted logical combinations without requiring them to be instructed in conventional Boolean logical terminology and representation.

The keyword matching is preferably implemented with a disregard for case (so that "SYSTEM", "System", and "system" will not be distinguished); and with the implicit condition that the required keyword is not followed by a letter character (so that "AI" does not match "main", but "system" matches "system" (with a final space) and "system." (with a final full stop), and also matches "subsystem"). Obviously, variations and elaborations of these conventions are possible; for example, a plural formed by adding a final "s" can be treated as equivalent to its singular. (The term "keyword" is not confined to a single word, but can include character combinations such as "A.I.".)

### Concluding remarks

The apparatus can be modified to allow one of the actions of a rule to be the total deletion of a message. This can be achieved by including an additional field in the actions part of the rule, or implicitly by treating the lowest priority level as a delete function. Deletion can of course be combined with forwarding of the message to someone else.

Another possible modification involves using arithmetical weighting to calculate the priority value of a message. For this, the subject matter of the subject field and/or text of the message (as determined by the keywords), the sender, the direct and copy-to addressees and their numbers are all given suitable weightings, and the sum of the weightings is quantized to give the priority value.

### Claims

1. In an electronic mail system comprising a plurality of stations between which messages can be sent, a message processing station comprising: a message store (11) for storing incoming messages;

a rules store (12) containing a plurality of rules relating to the status of received messages;  
 a rules test circuit (13) coupled to the message store (11) and to the rules store (12) and including comparison means (52) for comparing a message stored in the message store (11) with the rules to achieve a set of comparison results; and  
 routing means (14) for routing each message to a location selected by the set of comparison results for that message.

2. A message processing station according to Claim 1 in which the rules in the rules store (12) can be partitioned into a plurality of sets, and means (20) for selecting any desired one of said sets for use by the rules test circuit (13).

3. A message processing station as claimed in Claim 2 in which the locations selectable by the routing means include a folder store (15) including a plurality of folders.

4. A message station as claimed in Claim 3 in which the folder store further includes an in tray (16).

5. A message processing station as claimed in Claim 4 in which the rules test circuit further allocates priorities to at least the message stored in the in tray and stores them in the in tray in order of priority.

6. A message processing station as claimed in Claim 1 in which the locations selectable by the routing means include other stations of the system.

7. A message processing station as claimed in Claim 1 in which the rules store (12) further includes rule control apparatus (22) by means of which the rules can be added, amended, and deleted.

8. A message processing station as claimed in Claim 7 in which each set of rules is defined in relation to the needs of a particular user of the system including a keyboard (20) connected to the rules control apparatus (22) each rule being able to be amended or deleted by a user by operation of the keyboard (20).

9. A message processing station according to Claim 2 in which each rule comprises a tests part (35A) and an actions part (35B), the tests part including one or more of

a sender list field (37) to be matched against the sender (26) of a message,

an addressee limits values field (38) to be matched against the number of addressees and/or copy tos (28, 29) of the message,

an addressee list field (30) to be matched against the addressees (28) of the message.

a key-phrase field (40) to be matched against the body (31) of the message, and

a body-limits field (41), associated with the key-phrase field, which defines the extent of the body (31) of the message against which the key-phrase

field is matched,

and the actions part being operative on the message matching the tests part and including a priority value (45) for assigning to the message, and#

a location determining sub-part (46, 47) which determines the location to which the message is to be routed by the routing means (14).

10. A message processing station according to Claim 9 wherein the location determining sub-part of the actions part of the rule includes a file-to field (46) which determines which folder or folders (17, 18) of a folder store (15) the message is sent to.

11. A message processing station according to Claim 9 wherein the location determining sub-part of the actions part of the rule includes a forward-to field (47) which other stations of the system the message is to be sent to.

12. A method of controlling messages in an electronic mail system comprising a plurality of stations between which messages can be sent, the method including the steps of:

receiving and storing an incoming message (10, 11);

comparing the stored incoming message with a plurality of rules (13) to produce a set of comparison results for the message;

allocating at least one location (15) for subsequent storage of the message in accordance with the comparison results for that message.

13. A method of sorting messages in an electronic mail system as claimed in Claim 8 in which each message of a plurality of received messages is allocated an importance priority in accordance with the results of the comparison of each of the messages with the rules.

14. An electronic mail system comprising a plurality of stations between which messages can be sent, and including classification means in each station which assigns different priorities to the messages in dependence on matching of each message to a set of criteria defined by a set of rules, in which each station includes a message store for accumulating incoming messages and in which the classification means are adapted to operate on the messages so accumulated in the message store to effectively order the sequence in which the messages are presented to the user without changing the messages themselves.




Fig. 1






Fig. 2